

A SWITCHING CONCENTRATOR

CROSS-REFERENCE TO RELATED APPLICATION

This application is a non-provisional application of provisional application

5 Serial No. 60/212,333 filed June 16, 2000.

BACKGROUND OF THE DISCLOSURE

1. Field of the Invention

10 This invention relates to switching circuitry and, more particularly, to an m-to-n concentrator constructed from smaller concentrator/sorters or sorters.

2. Description of the Background Art

15 A switch with an array of m input ports and an array of n output ports is said to be a switch of “size” $m \times n$. One convention linearly labels the m input ports and the n output ports with, respectively, m and n distinct addresses. Usually, the addresses start from 0, so that the m input ports and n output ports are linearly labeled as 0, 1, ..., m-1 and 0, 1, ..., n-1, respectively (in a top-down manner if one visualizes a graphical representation of the input and output ports of the switch in a planar view). This inventive subject matter pertains to $m \times m$ switches and, particularly, to a special type of
20 switch designated a “concentrator”.

An $m \times m$ “sorter” is an $m \times m$ switch that compares the values of the m input signals and routes the m input signals to the m output ports such that the values of the m signals are monotonically increasing with respect to the output addresses. A number of

signals are said to be monotonically increasing with respect to the output addresses if the value of the signal routed to any particular output address is always not less than that of any other signal routed to a smaller output address. Similarly, a number of signals are said to be monotonically decreasing with respect to the output addresses if the value of the signal routed to any particular output address is always not less than that of any other signal routed to a larger output address. Here an input signal may possibly be an idle expression that is artificially generated at an input port when there is no real data input signal arriving at that input port. As per the graph representation, since the m output ports of an $m \times m$ switch are by default labeled from 0 to $m-1$ in the top-down manner, for a default $m \times m$ sorter, any signal routed to a lower output port cannot be less than those routed to upper output ports. Thus signal values on the output side are linearly sorted. For example, FIG. 1 shows a 6×6 sorter **100** as described above. The six input signals **101** assume different integer values and are linearly sorted on the output side **102** of the sorter according to their values.

In a wider sense, an $m \times m$ sorter can still be regarded as an $m \times m$ sorter even if the ordering of output ports is rotated, flipped, and/or even otherwise permuted in a predetermined fashion, as long as it observes the principle qualification of a sorter that the value of the signal routed to any particular output address is not less than that of any other signals routed to smaller output addresses. For example, FIG. 2 shows a 6×6 wide-sense sorter **200** which sorts the six input signals **201** into a rotation of an increasing sequence **202**. For simplicity, and without loss of generality, the sorters appearing in the sequel are of the default type as the one shown in FIG. 1 unless otherwise specified.

A 2×2 sorter is conventionally called a “sorting cell”. A default sorting cell always routes the larger one of its two local input signals to its lower output port. All sorting cells appearing hereafter are of the default version unless otherwise specified.

5 For $n < m$, an “m-to-n concentrator” is an $m \times m$ switch that routes the largest n among the m input signals to the n output ports with the n largest output addresses and the smallest $m-n$ among the m input signals to the $m-n$ output ports with the $m-n$ smallest output addresses. Thus the m output ports can be thought of being partitioned into a “1-output group” and a “0-output group”, with the former comprising

10 the n output ports with the n largest addresses and the latter comprising the remaining $m-n$ output ports. Then, an m-to-n concentrator can be regarded as a device which is capable of partitioning the m input signals (including real data input signals and artificial idle expressions) into two groups: the group of n largest signals, which are routed to the 1-output group, and the group of $m-n$ smallest signals, which are routed to the 0-output

15 group. As per the graph representation, by default the m-to-n concentrator is the one wherein the upper $m-n$ output ports form the 0-output group and the lower n output ports form the 1-output group. FIG. 3A shows a 6-to-2 concentrator **300**. The two signals of largest values **303** are routed to the 1-output group **301**, the group of the lower two output ports, wherein the order between these two signals is arbitrary. Meanwhile, the four

20 signals of smaller values **304** are routed to the 0-output group **302**, the group of the upper four output ports, within which the order among the signals is also arbitrary. Similar to the case of sorters, the definition of an m-to-n concentrator in a wider sense allows the rotation, flipping, and/or any other permutations, of the ordering of output ports.

Anyway, for any m-to-n concentrator, the 0-output group always comprises the m-n output ports with the smallest m-n output addresses while the 1-output group comprises the n output ports with the largest n output addresses. The only difference between the default and the wide-sense versions lies on the labeling of the m output addresses. For a default m-to-n concentrator, the m output ports are labeled from 0 to m-1 from top to bottom such that the 0-output group always comprises the upper m-n output ports and the 1-output group the lower n output ports. On the other hand, for a wide-sense m-to-n concentrator, the labeling of the m output ports, which is also from 0 to m-1, can be in any predetermined fashion, so the 0-output group can comprise any m-n distinct output ports and the 1-output group the remaining n output ports. FIG. 4 shows an example of a 6-to-2 concentrator **400** in such a wide sense wherein the six output ports are labeled as 0, 1, ..., 5 from bottom to top such that the two signals of largest values **403** are routed to the 1-output group **401**, which is now comprising the upper two output ports.

In some references in the background art, there is notion of an "m×n concentrator", which means an m×n switch, n<m, such that the largest n input signals are routed to the n output ports. Thus an m-to-n concentrator can be reduced to an "m×n concentrator" by not implementing output ports in the 0-output group. FIG. 3B shows a 6×2 concentrator **310** by which only the two largest signals are routed to the two output ports **311**. As stated above, this 6×2 concentrator may result from the 6-to-2 concentrator **300** in FIG. 3A with the upper four output ports **302** not implemented.

In order to avoid terminology ambiguity, the notion of an “ $m \times n$ concentrator” will not be adopted. Every concentrator in this context refers to an m -to- n concentrator for some m and some n , $n < m$.

5 For any $n < m$, by letting the lower n outputs of an $m \times m$ sorter form the 1-output group and the upper $m - n$ outputs form the 0-output group, the $m \times m$ sorter automatically qualifies as an m -to- n concentrator. One way to view the difference between an $m \times m$ sorter and a generic m -to- n concentrator is that the former provides a linear ranking among the n signals in the 1-output group and also a linear ranking among
10 the $m - n$ signals in the 0-output group.

For $n \leq m$, an “ m -to- n concentrator/sorter” is an $m \times m$ switch that routes the smallest $m - n$ among input signals to the $m - n$ smallest output addresses, which form the 0-output group, and the largest n among input signals to the n largest output addresses, which form the 1-output group, such that the values of these n signals are nondecreasing with respect to output addresses. In other words, an m -to- n concentrator/sorter not only is an m -to- n concentrator but also routes signals in a way such that the n signals reaching the 1-output group are sorted; the $m - n$ signals reaching the 0-output group, however, are not necessarily sorted. FIG. 5 shows a (default) 6-to-3 concentrator/sorter **500** as an
15 example wherein the three signals **501** of largest values are grouped together at the bottom and are also linearly sorted while the remaining smaller ones **502** are grouped at the top without sorting. Like the sorter and the concentrator, the concentrator/sorter also
20

has the wide-sense version, but the concentrator/sorters appeared in the context are of the default version unless otherwise specified.

For any $n \leq m$, by letting the lower n outputs of an $m \times m$ sorter form the 1-output group and the upper $m-n$ outputs form the 0-output group, the $m \times m$ sorter automatically qualifies as an m -to- n concentrator/sorter. Thus an $m \times m$ sorter is a special case of an m -to- n concentrator/sorter for all $n \leq m$.

Note that the comparison performed in a sorter, a concentrator, or a concentrator/sorter is based on the order defined over all of the possible values of an input signal, thus a sorter, a concentrator, or a concentrator/sorter should always be associated with an order. This order may vary from application to application and can even be any artificial one. For example, in certain applications of packet switching, the legitimate values of an input signal to a concentrator are "00", "10", and "11", and the concentrator is associated with an artificial order "10" < "00" < "11" which looks awkward but is practically very useful.

An m -to- n concentrator associated with a particular order among all possible values of an input signal automatically becomes an m -to- $(m-n)$ concentrator associated with the reverse of that order if the output ports of the m -to- n concentrator are re-labeled such that the n output ports in the 1-output group of the m -to- n concentrator, that is, the n output ports with the n largest addresses, are now labeled with the n smallest addresses and thus form the new 0-output group of the m -to- $(m-n)$ concentrator, and

similarly the 0-output group of the m-to-n concentrator now becomes the new 1-output group of the m-to-(m-n) concentrator. For example, consider FIG. 3A again. The 6-to-2 concentrator **300** associated with the natural order is equivalent to a 6-to-4 concentrator associated with the reversed natural order, that is, "0" > "1" > "2" > ..., when the output addresses are re-labeled in the reversed order. The same argument applies to the m-to-n concentrator/sorter.

In background art references, sorters, concentrators, concentrator/sorters, and other switching devices with sorting capabilities are collectively referred to as sorting devices. They can be constructed by various types of architectures, such as crossbar, shared-buffer memory, and multi-stage interconnection network of sorting cells. A multi-stage interconnection network means a collection of interconnected nodes, where the nodes are grouped into a number of stages such that every interconnection line is between two nodes on adjacent stages. Some architectures involve centralized control, which requires high processing and memory speeds and hence is suitable for implementing devices with small number of I/Os but inevitably imposes a bottleneck on the performance when the number of I/Os becomes large. Architectures in the type of multi-stage interconnection network of sorting cells can avoid the centralized control as follows. Every data unit includes an "in-band control signal" followed by a payload. When two data units respectively enter the two input ports of a sorting cell in the multi-stage interconnection network, the "in-band control signals" of the two data units are used as the "input signals to the sorting cell" for comparison. Note that this means the switching decision of the sorting cell is purely determined by just the two "in-band

control signals” of two data units and is independent of all other concurrent data units in the multi-stage interconnection network regardless the scale of the network. In this sense, the in-band control of a multi-stage network of sorting cells is extremely distributed.

5 When a data unit is routed through a multi-stage interconnection network, it may traverse a sorting cell on every stage of the network and its “in-band control signal” is preserved throughout. Since the switching control at all sorting cells on its route is by the simple comparison of the value of its “in-band control signal” against another value, it appears as if the routing of each individual signal through the network is guided by the
10 value of the signal itself. This distributed control mechanism is sometimes referred to as “self-routing” in the literature. Self-routing control over a multi-stage interconnection network of sorting cells enables the construction of large-scale switching devices.

There exist many ways to construct an m-to-n concentrator/sorter.

15 Algorithms for the construction of an m-to-n concentrator/sorter by multi-stage interconnection of sorting cells include, as representative of the art, the so-called “knockout tournaments” technique as disclosed by Y. S. Yeh, M. G. Hluchyj, and A. S. Acampora, “The Knockout Switch: A Simple Modular Architecture for High Performance Packet Switching,” IEEE J. Select. Areas Commun., vol. 5, pp. 1274-1283,
20 1987.

In a reference entitled "Concentrators in ATM switching," Comp. Sys. Sci. Eng., vol. 6, pp. 335-342, 1996, authored in S.-Y. R. Li and C.-M. Lau (Li-Lau), the authors devised and discussed a special m-to-n concentrator, where $m=2n$, as constructed from two $n \times n$ sorters and n sorting cells.

5

There are no teachings or suggestions of, and thus the art is devoid of, how to generalize the subject matter of the references, and especially of Li-Lau, to handle the construction of an m-to-n concentrator from two concentrator/sorters or sorters and n sorting cells where m is not necessarily equal to 2n.

10

SUMMARY OF THE INVENTION

15

Certain limitations and other shortcomings and deficiencies are obviated in accordance with the present invention by circuitry, and a concomitant methodology, for implementing an m-to-n concentrator from two concentrator/sorters or two sorters and n sorting cells in the case of m not necessarily equal to 2n.

20

In accordance with one broad system aspect of the present invention, an $m \times m$ switch having m input ports and m output ports, the switch arranged as an m-to-n concentrator, $n < m/2$, wherein m-n of the m output ports are grouped into a 0-output group and the remaining n output ports are grouped into a 1-output group, with the concentrator being composed of: (a) an $\lfloor m/2 \rfloor$ -to-n first concentrator/sorter wherein n of the $\lfloor m/2 \rfloor$ output ports are grouped into a 1-output group (where $\lfloor \cdot \rfloor$ is the conventional notation for "floor"); (b) an $\lceil m/2 \rceil$ -to-n second concentrator/sorter wherein n of the $\lceil m/2 \rceil$

output ports are grouped into a 1-output group (where $\lceil \cdot \rceil$ is the conventional notation for “ceiling”); and (c) n sorting cells wherein each of the sorting cells has a first input port connected to a specific one of the output ports of the 1-output group of the first concentrator/sorter and a second input port connected to a specific one of the output ports of the 1-output group of the second concentrator/sorter and wherein the n lower output ports of the sorting cells form the 1-output group for the concentrator.

In accordance with one method system aspect of the present invention, a method for concentrating the n largest of m incoming signals, $n < m/2$, includes: (a) processing $\lfloor m/2 \rfloor$ of the incoming signals with an $\lfloor m/2 \rfloor$ -to- n concentrator/sorter to produce a first set of comparison signals ordered in increasing order from 1 to n ; (b) processing the remaining $\lceil m/2 \rceil$ of the incoming signals with an $\lceil m/2 \rceil$ -to- n concentrator/sorter to produce a second set of comparison signals ordered in decreasing order from 1 to n ; and (c) selecting the larger between the k -th signal in the first set and the k -th signal of the second set, for each $k = 1, 2, \dots, n$, as one of the n largest input signals.

As noted in the above, an $m \times m$ sorter is a special case of an m -to- n concentrator/sorter for all $n \leq m$. Thus the $\lfloor m/2 \rfloor$ -to- n first concentrator/sorter in accordance with the present invention may be replaced by an $\lfloor m/2 \rfloor \times \lfloor m/2 \rfloor$ sorter. Similarly, the $\lceil m/2 \rceil$ -to- n second concentrator/sorter may be replaced by an $\lceil m/2 \rceil \times \lceil m/2 \rceil$ sorter. It is however worthwhile to note that, although an m -to- n concentrator can be constructed from two sorters and n sorting cells, normally a sorter is more complex than a

concentrator/sorter, so using concentrator/sorters as the building blocks as disclosed and claimed is generally more advantageous than using sorters.

Moreover, the inventive subject matter is independent of the construction algorithm for the two concentrator/sorters. In the special case when the two concentrator/sorters are constructed from multi-stage interconnection networks of sorting cells, the n sorting cells in the construction of a concentrator then naturally form an additional stage so that the whole concentrator is also a multi-stage interconnection network of sorting cells and thereby naturally suitable for the switching control in the self-routing fashion.

DESCRIPTIONS OF THE DRAWINGS

The teachings of the present invention can be readily understood by considering the following detailed description in conjunction with the accompanying drawings, in which:

FIG. 1 shows an example of $m \times m$ sorter;

FIG. 2 shows an example of $m \times m$ wide-sense sorter;

FIG. 3A shows an example of m -to- n concentrator;

FIG. 3B shows an example of $m \times n$ concentrator;

FIG. 4 shows an example of a wide-sense version of an m -to- n concentrator;

FIG. 5 shows an example of m -to- n concentrator/sorter;

FIG. 6 illustrates an example of a concentrator in accordance with the present invention for the case of two concentrator/sorters;

FIG. 7 illustrates an example of a concentrator in accordance with the present invention for the case of two sorters;

FIG. 8 illustrates the fundamental implementation of the m-to-n concentrator ($n < m/2$) in accordance with the present invention;

5 FIG. 9 is a flow diagram in accordance with a method of the present invention; and

FIG. 10 is an exemplary block diagram of two-stage interconnection of sorting cells.

To facilitate understanding, identical reference numerals have been used, 10 where possible, to designate identical elements that are common to the figures.

DETAILED DESCRIPTION

Consider the following examples to elucidate the teachings and suggestions of the present invention.

15 Example 1. With reference to FIG. 6, there is shown concentrator **600** for the case of $m=9$ and $n=3$, wherein concentrator **600** is composed of: (a) $\lfloor m/2 \rfloor$ -to- n concentrator/sorter **601**; (b) $\lceil m/2 \rceil$ -to- n concentrator/sorter **610**; and (c) n sorting cells **602-1**, **602-2**, and **602-3** interconnected to elements **601** and **610** in the following manner. Each of the sorting cells receives one signal from concentrator/sorter **601** and one signal 20 from concentrator/sorter **610** in such a way that: sorting cell **602-1** compares the largest signal from the 1-output group **604** of element **610** with the smallest signal the 1-output group **605** of element **601** and routes the larger one to the 1-output group of the concentrator, that is, the group of n output ports **603**; sorting cell **602-2** compares the

second largest signal from the 1-output group of element **610** with the second smallest signal from the 1-output group of element **601** and also routes the larger one to the group of n output ports **603**; and, finally, sorting cell **602-3** compares the third largest signal from the 1-output group of element **610** with the third smallest signal from the 1-output group of element **601** and also routes the larger one to the group of n output ports **603**.

Note that some 1×1 delay elements, like the elements **620-1**, **620-2** and **620-3**, are sometimes inserted to ensure the synchronization of the signals. The delay elements are usually not explicitly shown in the graph representation when there is no ambiguity.

Consider the specific construction identified in FIG. 6; in particular, **601** is a 4-to-3 concentrator/sorter in the default version while **610** is a 5-to-3 wide-sense concentrator/sorter whose five output ports are linearly addressed as 0, 1, ..., 4 from bottom to top such that the 1-output group comprises the upper three output ports. Four input signals having values 1, 6, 4, and 10 appear in top-down manner as inputs to element **601**, and five input signals having values 9, 8, 10, 1, and 2 appear in top-down manner as inputs to element **610**. Element **601**, by its 4-to-3 concentrator/sorter functionality, produces output signals 1, 4, 6, 10 read from the top to the bottom.

Similarly, element **610** produces output signals 10, 9, 8, 1, and 2 in top-down manner.

Recall that a concentrator/sorter only linearly orders the n largest output signals, and the other $n-m$ output signals need not be linearly ordered. For element **601**, the $n=3$ largest output signals of 4, 6, and 10 are linearly ordered; for element **610** output signals 10, 9, and 8 are linearly ordered. The smallest two output signals from element **610** are in arbitrary order.

When the largest output signal from the 1-output group of element **610** is compared to the smallest output signal from the 1-output group of element **601**, that is, when the signal with the value of “4” is compared to the signal with the value of “10”, sorting cell **602-1** produces “4” on its upper lead and “10” on its lower lead which leads to the bottom output port of the 1-output group, that is, the group of n output ports identified by reference numeral **603**. Similarly, sorting cell **602-2** produces “9” at the next-to-bottom output port of group **603**, and sorting cell **602-3** produces “10” at the remaining output port of group **603**. Notice that the three signals at the output ports of group **603** are the largest three among the input signals, but are not ordered as per the functionality of a concentrator.

Accordingly, the 1-output group of the concentrator includes the lower output port of each of the n sorting cells, while the 0-output group of the concentrator includes the upper output port of each of the n sorting cells plus the $\lfloor m/2 \rfloor - n$ output ports of the 0-output group of the first concentrator/sorter and the $\lceil m/2 \rceil - n$ output ports of the 0-output group of the second concentrator/sorter.

Example 2. With reference to FIG. 7, there is shown concentrator **700** for the case of $m=9$ and $n=3$, wherein concentrator **700** is composed of: (a) $\lfloor m/2 \rfloor \times \lfloor m/2 \rfloor$ sorter **701**; (b) $\lceil m/2 \rceil \times \lceil m/2 \rceil$ sorter **710**; and (c) n sorting cells **702-1**, **702-2**, and **702-3** interconnected to elements **701** and **710** in the following manner. Each of the sorting cells receives one signal from sorter **701** and one signal from sorter **710** in such a way that: sorting cell **702-1** compares the largest signal from element **710** with the n^{th} largest

signal, or equivalently, the smallest signal out of the n largest signals, from element **701** and routes the larger one to the 1-output group of the concentrator, that is, the group of n output ports **703**; sorting cell **702-2** compares the second largest signal from element **710** with the $(n-1)^{\text{th}}$ largest signal, or equivalently, the second smallest signal out of the n largest signals, from element **701** and also routes the larger one to the group of n output ports **703**; and, finally, sorting cell **702-3** compares the third largest signal from element **710** with the third smallest signal out of the n largest signals, that is, the largest signal from element **701** and also routes the larger one to the group of n output ports **703**.

Consider the specific construction identified in FIG. 7; in particular, **701** is a 4×4 sorter in the default version while **710** is a 5×5 wide-sense sorter whose five output ports are linearly addressed as 0, 1, ..., 4 from bottom to top. Four input signals having values 1, 6, 4, and 10 appear in top-down manner as inputs to element **701**, and five input signals having values 9, 8, 10, 1, and 2 appear in top-down manner as inputs to element **710**. Element **701**, by its sorter functionality, produces output signals 1, 4, 6, 10 read from the top to the bottom. Similarly, element **710** produces output signals 10, 9, 8, 1, and 2 in top-down manner. Recall that a sorter produces linearly ordered output signals, which is readily apparent from the output signals of both elements **701** and **710**. For element **701**, the $n=3$ largest output signals are 4, 6, and 10; for element **710** the three largest output signals are 10, 9, and 8.

When the largest output signal from element **710** is compared to the smallest output signal from element **701**, that is, when the signal with the value of “4” is compared to the signal with the value of “10”, sorting cell **702-1** produces “4” on its

upper lead and “10” on its lower lead which leads to the bottom output port of the 1-output group, that is, the group of n output ports identified by reference numeral **703**. Similarly, sorting cell **702-2** produces “9” at the next-to-bottom output port of group **703**, and sorting cell **702-3** produces “10” at the remaining output port of group **703**. Notice

5 that the three signals at the output ports of group **703** are the largest three of the input signals, but are not ordered as per the functionality of a concentrator.

Generic Implementation Using Two Concentrator/Sorters

As noted in the above, an m -to- n concentrator associated with a particular order

10 among all possible values of an input signal is automatically an m -to- $(m-n)$ concentrator associated with the reverse of that order, although the m -to- $(m-n)$ concentrator is now in the wide sense with its output addresses re-labeled such that the n output ports forming the 1-output group of the m -to- n concentrator are now labeled with the n smallest

15 addresses and hence form the 0-output group of the m -to- $(m-n)$ concentrator while the 0-output group of the m -to- n concentrator becomes the 1-output group of the m -to- $(m-n)$ concentrator. Thus any methodology for m -to- n concentration can also be employed in m -to- $(m-n)$ concentration. Because of this symmetry, to derive a methodology for m -to- n concentration, one may derive a methodology for m -to- $(m-n)$ concentration instead, or

20 employ an existing methodology already developed for m -to- $(m-n)$ concentration. Since either $n \leq m/2$ or $m-n \leq m/2$ holds, it suffices to consider m -to- n concentration with $n \leq m/2$ hereafter.

A fundamental implementation of the present invention is depicted in FIG.

8, which is a generalization of FIG. 6. Concentrator **800** includes element **801** which represents an $\lfloor m/2 \rfloor$ -to- n concentrator/sorter wherein the n output addresses in its 1-output group increases from top to bottom and thus the n largest signals are linearly sorted into a monotonically increasing sequence from top to bottom (the arrow at the output points towards the largest elements); signal leads shown by reference numeral **803** is thus a group of $\lfloor m/2 \rfloor$ external input signals while the signal leads shown by reference numeral **805** depicts the 1-output group of the concentrator/sorter **801**. Similarly, element **810** represents an $\lceil m/2 \rceil$ -to- n concentrator/sorter wherein the n output addresses in its 1-output group decreases from top to bottom and thus the n largest signals are linearly sorted into a monotonically decreasing sequence from top to bottom (the arrow at the output points towards the largest elements); reference numeral **804** refers to a group of $\lceil m/2 \rceil$ external input signals while reference numeral **806** refers to the 1-output group of the concentrator/sorter **810**. (Alternatively, as is readily appreciated from Example 2 and FIG.7, element **801** may stand for an $\lfloor m/2 \rfloor \times \lfloor m/2 \rfloor$ sorter and element **810** an $\lceil m/2 \rceil \times \lceil m/2 \rceil$ sorter with their output lists respectively ordered to achieve the same outputs previously described.) Each of the elements denoted by reference numerals **802-1**, **802-2**, ..., **802-n** is a sorting cell, i.e. the larger of the two input signals is routed to the lower output port and the smaller to the upper output port. There are n parallel sorting cells **802**. Each of these sorting cells receives one signal from the group **805** and one from the group **806** in such a way that one of the sorting cells **802-1** compares the largest signal from **806** with the smallest signal from **805** and routes the larger one to the 1-output group **807** of the m -to- n concentrator, another one (**802-2**) compares the second

largest signal from **806** with the second smallest signal from **805** and also route the larger one to the group **807**, and so on. This generates as the 1-output group **807** the n largest signals among the $2n$ signals in groups **805** and **806** altogether, which also are the largest n signals of the m signals in groups **803** and **804** altogether. This completes the desired concentration action.

To be more concise, if the n output signals in the 1-output group **805** of first concentrator/sorter **801** are ordered from 1 to n in association with increasing signal values and the n output signals in the 1-output group **806** of second concentrator/sorter **810** are ordered from 1 to n in association with decreasing signal values, then n sorting cells **802-1**, ..., **802-n** are interconnected with the two concentrator/sorters such that the k -th output signal in the 1-output group of first concentrator/sorter **801** and the k -th output signal in the 1-output group of second concentrator/sorter **810** serve as the input signals to the k -th sorting cell (that is, sorting cell **802-k**) for $k = 1, 2, \dots, n$.

It is noted that in FIG. 8, element **801** can be $\lfloor m/2 \rfloor \times \lfloor m/2 \rfloor$ and element **810** can be $\lceil m/2 \rceil \times \lceil m/2 \rceil$, which two combine to give a group of m inputs and a group of m outputs.

Element **801** can, of course, also be $\lceil m/2 \rceil \times \lceil m/2 \rceil$ and element **810** be $\lfloor m/2 \rfloor \times \lfloor m/2 \rfloor$.

It is further readily appreciated that the advantage to implementing concentrators as illustrated by the arrangement of FIG. 8 is that smaller building-block components are combined to realize a larger concentrator. In the special case when the two concentrator/sorters are constructed from multi-stage interconnection networks of sorting cells, the concentrator so constructed by the teachings of FIG. 8 is then also a multi-stage interconnection network of sorting cells and thereby naturally suitable for the switching control in the self-routing fashion known in the prior art.

Flow Diagram

Flow diagram **900** of FIG. 9 depicts the methodology for m-to-n concentration, $n < m/2$. Processing block **910** effects the operation of processing $\lfloor m/2 \rfloor$ of the incoming signals with an $\lfloor m/2 \rfloor$ -to-n concentrator/sorter to produce a first set of comparison signals ordered in increasing order from 1 to n. In turn, processing block **920** invokes processing of the remaining $\lceil m/2 \rceil$ of the incoming signals with an $\lceil m/2 \rceil$ -to-n concentrator/sorter to produce a second set of comparison signals ordered in decreasing order from 1 to n. Finally, processing block **930** is used to select the larger between the k-th signal of the first set and the k-th signal of the second set, for each $k = 1, 2, \dots, n$, as one of the n largest input signals.

From another viewpoint, the concentrator as described above, such as in FIG. 8, is an $m \times m$ switch that routes m incoming signals in a special manner that achieves m-to-n concentration. The concomitant method underlying such a switch may then be described as a method for routing the n largest of m incoming signals, $n < m/2$, to n outputs, wherein the method includes: (a) processing $\lfloor m/2 \rfloor$ of the incoming signals

with an $\lfloor m/2 \rfloor$ -to- n concentrator/sorter to produce a first set of comparison signals ordered in increasing order from 1 to n ; (b) processing the remaining $\lceil m/2 \rceil$ of the incoming signals with an $\lceil m/2 \rceil$ -to- n concentrator/sorter to produce a second set of comparison signals ordered in decreasing order from 1 to n ; and (c) routing the larger

5 between the k -th signal in the first set and the k -th signal of the second set to one of the n outputs, $k = 1, 2, \dots, n$.

Multistage Interconnection Network of Sorting Cells

The block diagram of FIG. 10 shows a two-stage interconnection network

10 of sorting cells which is representative of a generic multi-stage interconnection network of sorting cells. In particular, FIG. 10 depicts a 4×4 2-stage network 1010 wherein each node 1021, 1022, 1023, and 1024 is filled with a sorting cell. The external input addresses (0, 1, 2, 3) are in natural order, and their correspondence with input addresses of input nodes **1021**, **1022** is readily discerned. Similarly the external output addresses

15 (0, 1, 2, 3) are in natural order and their correspondence with output addresses of output nodes **1023**, **1024** is readily discerned.

Although the present invention has been shown and described in detail herein, those skilled in the art can readily devise many other varied embodiments that still

20 incorporate these teachings. Thus, the previous description merely illustrates the principles of the invention. It will thus be appreciated that those with ordinary skill in the art will be able to devise various arrangements which, although not explicitly described or shown herein, embody principles of the invention and are included within its spirit and

scope. Furthermore, all examples and conditional language recited herein are principally intended expressly to be only for pedagogical purposes to aid the reader in understanding the principles of the invention and the concepts contributed by the inventor to furthering the art, and are to be construed as being without limitation to such specifically recited

5 examples and conditions. Moreover, all statements herein reciting principles, aspects, and embodiments of the invention, as well as specific examples thereof, are intended to encompass both structural and functional equivalents thereof. Additionally, it is intended that such equivalents include both currently known equivalents as well as equivalents developed in the future, that is, any elements developed that perform the function,

10 regardless of structure.

In addition, it will be appreciated by those with ordinary skill in the art that the block diagrams herein represent conceptual views of illustrative circuitry embodying the principles of the invention.

15